

On the Topologies Formed by Selfish Peers

Thomas Moscibroda, Stefan Schmid, Roger Wattenhofer
{moscitho, schmiste, wattenhofer}@tik.ee.ethz.ch
Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland

ABSTRACT

Current peer-to-peer (P2P) systems often suffer from a large fraction of freeriders not contributing any resources to the network. Various mechanisms have been designed to overcome this problem. However, the selfish behavior of peers has aspects which go beyond resource sharing. This paper studies the effects on the topology of a P2P network if peers selfishly select the peers to connect to. In our model, a peer exploits locality properties in order to minimize the latency (or response times) of its lookup operations. At the same time, the peer aims at not having to maintain links to too many other peers in the system. We show that the resulting topologies can be much worse than if peers collaborated. Moreover, the network may never stabilize, even in the absence of churn.

1. INTRODUCTION

The power of peer-to-peer (P2P) computing arises from the collaboration of its numerous constituent parts, the peers. If all the participating peers contribute some of their resources—for instance bandwidth, memory, or CPU cycles—, highly scalable decentralized systems can be built which significantly outperform existing server based solutions. Unfortunately, in reality, many peers are selfish and strive for maximizing their own utility by benefiting from the system without contributing much themselves. Hence the performance—and thus its success in practice!—of a P2P system crucially depends on its capability of dealing with *selfishness*. A well-known mechanism designed to cope with this freeriding problem is the *tit-for-tat policy*, as for instance employed by the file-distribution tool *BitTorrent*.

However, selfish behavior in peer-to-peer networks has numerous important implications even beyond the peer's unwillingness to contribute bandwidth or memory. For example, in unstructured P2P systems—the predominant P2P architectures in today's Internet—, a peer can select to which and to how many other peers in the network it wants to connect. With a clever choice of neighbors, a peer can attempt to optimize its lookup performance by minimizing the latencies—or more precisely, the *stretch*—to the other peers in the network. Achieving good stretches by itself is of course simple: A peer can establish links to a large number of other peers in the system. Because the memory and maintenance overhead of such a neighbor set is large, however, egoistic peers try to exploit locality as much as possible, while avoiding to store too many neighbors. It is this fundamental trade-off between the need to have small latencies and the desire to reduce maintenance overhead that governs the decisions of selfish peers.

This paper investigates the impact of selfish neighbor selection on the quality of the resulting network topologies. An appropriate tool to study such selfish behavior is *game theory*. In particular, this paper studies the *Price of Anarchy* of P2P overlay creation, which is the ratio between an optimal solution obtained by perfectly collaborating participants compared to a solution generated by peers that act in an egoistic manner, optimizing their individual benefit. The importance of studying the Price of Anarchy in peer-to-peer

systems stems from the fact that it quantifies the possible degradation caused by selfishness. Specifically, a low Price of Anarchy indicates that a system does not require an incentive-mechanism (such as *tit-for-tat*), because selfishness does not overly bog down the overall system performance. If the Price of Anarchy is high, however, specific cooperation incentives (whose goal is to reduce the Price of Anarchy) need to be enforced in order to ensure that the system can perform efficiently. In peer-to-peer systems therefore, the Price of Anarchy is a measure that helps explaining the necessity (or non-necessity) of cooperation mechanisms in various aspects of these systems.

The contribution of this paper is twofold. First, we show that the topologies of selfish, unstructured P2P systems can be much worse than in a scenario in which peers collaborate. More precisely, we show that the Price of Anarchy is $\Theta(\min(\alpha, n))$, where α is a parameter that captures the tradeoff between lookup performance (low stretches) and the cost of neighbor maintenance, and n is the number of peers in the system, respectively. Thereby, the upper bound $O(\min(\alpha, n))$ holds for peers located in *arbitrary* metric spaces, including the popular *growth-bounded* and *doubling metrics*. On the other hand, intriguingly, this bound is tight even in such a simple metric space as the 1-dimensional *Euclidean space*. As a second contribution, we prove that the topology of a static peer-to-peer system consisting of selfish peers may never converge to a stable state. That is, links may continuously change even in environments without *churn* (causing the network to be inherently instable).

The remainder of the paper is organized as follows. In Section 2, we present our model and introduce the game theoretic approach. Related work is reviewed in Section 3. In the subsequent Section 4 we give tight bounds on the Price of Anarchy of P2P topologies. Then, in Section 5, we show that a system consisting of selfish peers may never stabilize, even if there is no churn. Finally, Section 6 concludes the paper.

2. MODEL

We model the peers of a P2P network as points in a *metric space* $\mathcal{M} = (V, d)$, where $d : V \times V \rightarrow [0, \infty)$ is the *distance function* which describes the underlying latencies between all pairs of peers.

The effects of selfish peer behavior is studied from a *game-theoretic* perspective. We consider a set of n peers

$$V = \{\pi_0, \pi_1, \dots, \pi_{n-1}\}.$$

A peer can choose to which subset of other peers it wants to store pointers (IP addresses). Formally, the *strategy space* of a peer π_i is given by $S_i = 2^{V \setminus \{\pi_i\}}$, and we will refer to the actually chosen links as π_i 's *strategy* $s_i \in S_i$. We say that π_i *maintains or establishes a link to* π_j if $\pi_j \in s_i$. The combination of all peers' strategies, i.e., $s = (s_0, \dots, s_{n-1}) \in S_0 \times \dots \times S_{n-1}$, yields a (directed) graph $G[s] = (V, \cup_{i=0}^{n-1} (\{\pi_i\} \times s_i))$, which describes the resulting P2P topology.

Selfish peers exploit *locality* in order to maximize their lookup

performance. Concretely, a peer aims at minimizing the *stretch* to all other peers. The stretch between two peers π and π' is defined as the shortest distance between π and π' using the links of the resulting P2P topology G divided by the direct distance, i.e., for a topology G , $stretch_G(\pi, \pi') = d_G(\pi, \pi')/d(\pi, \pi')$. Clearly, it is desirable for a peer to have low stretch to other peers in order to keep its latency small. By establishing a link to all peers in the system, a peer reaches every peer with minimal stretch 1, and the potential lookup performance is optimal. However, storing and especially maintaining a large number of links is expensive.¹ Hence, the individual cost $c_i(s)$ incurred at a peer π is composed not only of the stretches to all other peers, but also of its *degree*, i.e., the number of its neighbors:

$$c_i(s) = \alpha \cdot |s_i| + \sum_{i \neq j} stretch_{G[s]}(\pi_i, \pi_j).$$

Note that this cost function captures the classic P2P trade-off between the need to minimize latencies and the desire to store and maintain only few links, as it has been addressed by many existing systems, for example Pastry [11]. Thereby, the relative importance of degree costs versus stretch costs is expressed by the parameter α .

The objective of a selfish peer is to minimize its individual cost. In order to evaluate the topologies constructed by such selfish peers—and compare them with the topologies achieved by collaborating peers—we use the notion of a *Nash equilibrium*. A P2P topology constitutes a Nash equilibrium if no peer can reduce its individual cost by changing its set of neighbors given that the connections of all other peers remain the same. More formally, a (pure) Nash equilibrium is a combination of strategies s such that, for each peer π_i , and for all alternative strategies s' which differ only in the i^{th} component (different neighbor sets for peer π_i), $c_i(s) \leq c_i(s')$. This means that in a Nash equilibrium, no peer has an incentive to change its current set of neighbors, that is, Nash equilibria are *stable*.

While peers try to minimize their individual cost, the system designer is interested in a good overall quality of the P2P network. The *social cost* is the sum of all peers' individual costs, i.e.,

$$C(G) = \sum_i c_i = \alpha|E| + \sum_{i \neq j} stretch_G(\pi_i, \pi_j).$$

The lower this social cost, the better is the system's performance.

Determining the parameter α in real unstructured peer-to-peer networks is an interesting field for study. As mentioned, α measures the relative importance of low stretches compared to the peers' degrees, and thus depends on the system or application: For example, in systems with many lookups where good response times are crucial, α is smaller than in distributed archival storage systems consisting mainly of large files.² In the sequel, we denote link and stretch costs by $C_E(G) = \alpha|E|$ and $C_S(G) = \sum_{i \neq j} stretch_G(\pi_i, \pi_j)$, respectively.

Typically, a given distribution of peers in a metric space can result in different Nash equilibria, depending on the order in which peers change their links. To gain an understanding of the impact of selfishness on the social cost, we are particularly interested in the social cost of the *worst* possible Nash equilibrium. That is, we study topologies in which no selfish peer has an incentive to change its neighbors, but in which all peers together could be much better

¹For instance, the maintenance of a link may involve periodic pings to verify whether the neighbor is still alive.

²If α is in the order of $\Theta(\sqrt{n})$, for instance, P2P topologies in which the latency stretch between all pairs of peers is bounded by a constant, and in which every peer has at most degree $O(\sqrt{n})$ can be shown to be asymptotically optimal. This trade-off has for example been achieved by the Tulip system proposed in [1].

off if they collaborated. More precisely and using the terminology of game theory, we are interested in the *Price of Anarchy*, the ratio between the social cost of the worst Nash equilibrium and the social cost of the optimal topology.

3. RELATED WORK

The lack of cooperation in traditional P2P file-sharing systems has been well-documented over the last years [3, 13], and research on the causes and possible counter-measures is very active, e.g., [5] and [9]. Most of the current literature focuses on the issue of free resource consumption, *freeriding*. In contrast, the impact of other aspects of selfishness has received much less attention. In fact, to the best of our knowledge, this is the first paper to take a step towards studying the consequences of selfish neighbor selection on the topologies of P2P networks.

The first paper to study the creation of networks from a game-theoretic point of view is due to Fabrikant et al. [8]. In this paper, the authors analyze the Internet's architecture as built by economic agents, e.g., by Internet providers or *autonomous systems*. Recent subsequent work on network creation in various settings includes [4, 6, 7]. In contrast to all these works, our model takes into account many of the intrinsic properties of P2P systems. For instance, it captures the important *locality properties* of P2P systems, i.e., the desire to reduce the latencies (expressed as the stretch) experienced when performing look-up operations. Furthermore, the fact that a peer can decide to which other peers it wishes to store pointers yields a scenario with *directed* links.

Building structured systems that explicitly exploit locality properties has been a flourishing research area in networking and P2P computing (e.g. [1, 11, 12]). In early literature on distributed hash tables (DHT), the major measure of system quality has been the number of hops required for look-up operations. While this hop-distance is certainly of importance, it has been argued that the delay of communication (i.e., the stretch between pairs of peers) is a more relevant quality measure. Based on results achieved in [10], systems such as [1, 2, 11, 14] guarantee a provably bounded stretch with a limited number of links per peer. All of these systems are structured and peers are supposed to participate in a carefully predefined topology. Our paper complements this line of research by analyzing topologies as they are created by *selfish peers*, which are interested only in optimizing their *individual* trade-off between locality and maintenance overhead.

4. PRICE OF ANARCHY

The Price of Anarchy is a measure to bound the degradation of a globally optimal solution caused by selfish individuals. In this section, we show that the topologies created by selfish peers deteriorate more (compared to collaborative networks) as the cost of maintaining links becomes more important (larger α). Concretely, in Section 4.1 we prove that for *arbitrary* metrics—thus, including the important and well-studied *growth-bounded* and *doubling metrics*—, the Price of Anarchy never exceeds $O(\min(\alpha, n))$. We then show in Section 4.2 that this bound is tight even in the “simplest” metric space, the 1-dimensional Euclidean space.

4.1 Upper Bound

Assume the most general setting where n peers are arbitrarily located in a given metric space \mathcal{M} , and consider a peer π which has to find a suitable neighbor set. Clearly, the *maximal* stretch from π to any other peer π' in the system is at most $\alpha + 1$: If $stretch(\pi, \pi') > \alpha + 1$, π could establish a direct link to π' , reducing the stretch from more than $\alpha + 1$ to 1, while incurring




Figure 1: Example topology G where the Price of Anarchy is $\Theta(\min(\alpha, n))$ for $3.4 \leq \alpha$.

a link cost of α . Therefore, in any Nash equilibrium, no stretch exceeds $\alpha + 1$. Because there are at most $n(n-1)$ directed links (from each peer to all remaining peers), the social cost of a Nash equilibrium is $O(\alpha n^2)$. Since the optimum social costs is clearly lower bounded by $\Omega(\alpha n + n^2)$, we have the following result.

THEOREM 4.1. *For any metric space \mathcal{M} , the Price of Anarchy is $O(\min(\alpha, n))$.*

4.2 Lower Bound

We now show that there are P2P networks in with a Price of Anarchy of $\Omega(\min(\alpha, n))$, which implies that the upper bound of Section 4.1 is asymptotically tight. Intriguingly, the Price of Anarchy can deteriorate to $\Theta(\min(\alpha, n))$ even if the underlying latency metric describes a simple 1-dimensional Euclidean space.

Consider the topology G in Figure 1 in which peers are located on a 1-dimensional Euclidean line, and the distance (latency) between two consecutive peers increases exponentially towards the right. Concretely, peer i is located at position $\alpha^{i-1}/2$ if i is odd, and at position α^{i-1} if i is even. The peers of G maintain links as follows: All peers have a link to their nearest neighbor on the left. Odd peers additionally have a link to the second nearest peer on their right. In the following, we prove that G constitutes a Nash equilibrium. Afterwards, we derive the lower bound on the Price of Anarchy by computing the social cost of this topology.

LEMMA 4.2. *The topology G shown in Figure 1 forms a Nash equilibrium for $\alpha \geq 3.4$.*

PROOF. In the following, a proof sketch is given only.

We distinguish between even and odd peers. For both cases, we show that no peer has an incentive to deviate from its strategy.

Case even peers: Every even peer i needs to link to at least one peer on its left, otherwise i cannot reach the peers $j < i$. A connection to peer $i-1$ is optimal, as the stretch to all peers $j < i$ becomes 1. Observe that every alternative link to the left would imply a larger stretch to at least one peer on the left without reducing the stretch to peers on the right. Furthermore, i cannot reduce the distance to any—neither left nor right—peer by adding further links to the left. Hence, it only remains to show that i cannot benefit from adding more links to the right.

By adding a link to the right, peer i shortens the distance to *all* peers on the right. However, the cost reduction per peer decreases as a geometric series, and any such link to the right would strictly increase i 's costs. To show this, we consider two cases in turn: i linking to an odd peer on the right, and i linking to an even peer on the right.

Link to an odd peer: Consider the benefit of i adding a link to its odd neighbor $i+1$. For an odd peer $j > i$, we define the *benefit* $B_{i,j}$ as the stretch cost reduction caused by the addition of the link $(i, i+1)$. We have, for $i \geq 2$,

$$\begin{aligned} B_{i,j} &= \text{stretch}_{\text{old}}(i, j) - \text{stretch}_{\text{new}}(i, j) \\ &= \frac{d(i, i-1) + d(i-1, j)}{d(i, j)} - \frac{d(i, j)}{d(i, j)} \\ &= \frac{2 - 1/\alpha}{1/2\alpha^{j-i} - 1}. \end{aligned}$$

Similarly, the savings $B_{i,j}$ for an even peer $j > i$ and $i \geq 2$ amount to $B_{i,j} = (2 - 1/\alpha)/(\alpha^{j-i} - 1)$. Hence, for all $\alpha \geq 3.4$, the total savings B_i for peer i are less than

$$\begin{aligned} B_i &= \sum_{\text{odd } j > i} B_{i,j} + \sum_{\text{even } j > i} B_{i,j} \\ &< \sum_{\delta=1}^{\infty} \frac{2 - \frac{1}{\alpha}}{\frac{1}{2}\alpha^{2\delta-1} - 1} + \sum_{\delta=1}^{\infty} \frac{2 - \frac{1}{\alpha}}{\alpha^{2\delta} - 1} \\ &< \frac{4\alpha^2 - 1}{\alpha^2 - 1} \underset{(\alpha \geq 3.4)}{<} \alpha + 1. \end{aligned}$$

Therefore, the construction of link $(i, i+1)$ would be of no avail (benefit smaller than cost). Clearly, the benefit of alternative or additional links to odd neighbors on the right is even smaller.

Link to an even peer: A link to an even peer $j > i$ entails a stretch 1 to the corresponding peer instead of $\text{stretch}_{\text{old}}(i, j) = (\alpha^j - \alpha^{j-1} + \alpha^{i-1} - \alpha^{i-2})/(\alpha^{j-1} - \alpha^{i-1}) < \alpha + 1$ for $\alpha > 2$. However, the stretch from i to all other peers remains unchanged, since the path $i \rightsquigarrow (i-1) \rightsquigarrow (i+1)$ is shorter than $i \rightsquigarrow (i+2) \rightsquigarrow (i+1)$: $\alpha^{i-1} - \alpha^{i-2}/2 + \alpha^i/2 - \alpha^{i-2}/2 < \alpha^{i+1} - \alpha^{i-1} + \alpha^{i+1} - \alpha^i/2$ for $\alpha > 1$. Therefore, an even peer i has no incentive to build links to any even peer on its right.

Case odd peers: The proof that an odd peer i has no incentive to change its neighbor set is similar to the proof for even peers, and it is omitted here.

Having verified that the topology of Figure 1 is a Nash equilibrium, its social cost can be computed.

LEMMA 4.3. *The social cost $C(G)$ of the topology G shown in Figure 1 is $C(G) \in \Theta(\alpha n^2)$.*

PROOF. Clearly, the link costs of topology G are $C_E(G) \in \Theta(\alpha n)$. But since the stretch from an odd peer i to an even peer $j > i$ and the stretch between two even peers i and $j > i$ are $\text{stretch}(i, j) > \alpha/2$ (for $\alpha > 2$), the stretch costs are $C_S \in \Theta(\alpha n^2)$.

THEOREM 4.4. *The Price of Anarchy of the peer topology G shown in Figure 1 is $\Theta(\min(\alpha, n))$.*

PROOF. The theorem follows from Theorem 4.1 and Lemmas 4.2 and 4.3, and by the observation that the optimal social cost of a topology connecting the peers in Figure 1 is upper bounded by $O(\alpha n + n^2)$. For the latter, assume that there are no links in Figure 1. If every peer connects to the nearest peer on its left and to the nearest peer on its right, there are $2(n-1)$ links, and all stretches are 1. Thus, the social cost of this resulting topology \tilde{G} is $C(\tilde{G}) = \alpha \cdot 2(n-1) + n(n-1) \in O(\alpha n + n^2)$. The optimal social cost is at most the social cost of \tilde{G} .

5. EXISTENCE OF NASH EQUILIBRIA

In this section, we show that a system of selfish peers may never converge to a stable state, even in the absence of churn, mobility, or other sources of dynamism. Interestingly, this result even holds if we assume latencies to form simple metric spaces, such as a 2-dimensional Euclidean space.

THEOREM 5.1. *Regardless of the magnitude of α , there are metric spaces \mathcal{M} , for which there exists no pure Nash equilibrium, i.e. certain P2P networks cannot converge to a stable state. This is the case even if \mathcal{M} is a 2-dimensional Euclidean space.*

Instead of presenting the formal proof, we attempt to highlight the main ideas only. Assume that the parameter α is a multiple




Figure 2: Instance I_k has no pure Nash equilibrium when $\alpha = 0.6k$, where $k = n/5$. The number of peers in each cluster is k .

of 0.6, i.e., $\alpha_k = 0.6k$ for an arbitrary integer $k > 0$. Given a specific k , we show that the 2-dimensional Euclidean instance I_k of Figure 2 has no pure Nash equilibrium. Specifically, I_k constitutes a situation in which there are peers $\pi_1 \in \Pi_1$ and $\pi_2 \in \Pi_2$ that continue to deviate to a better strategy ad infinitum, i.e., the system cannot converge.

The n peers of instance I_k are grouped into five clusters $\Pi_1, \Pi_2, \Pi_a, \Pi_b$, and Π_c , each containing $k = n/5$ peers. Within a cluster, peers are located equidistantly on a line, and each cluster's diameter is ϵ/n , where $\epsilon > 0$ is an arbitrarily small constant. The *inter-cluster distance* $d(\Pi_i, \Pi_j)$ between Π_i and Π_j is the minimal distance between any two peers in the two clusters. Distances not explicitly defined in Figure 2 follow implicitly from the constraints imposed by the underlying Euclidean plane. A link from a peer $\pi_i \in \Pi_i$ to a peer $\pi_j \in \Pi_j$ is denoted by ℓ_{ij} . Clusters Π_a, Π_b , and Π_c are called *top-clusters*, Π_1 and Π_2 are *bottom-clusters*, and finally, δ denotes an arbitrarily small positive number such that $\delta > 10\epsilon$.

The proof unfolds in a series of lemmas that characterize the structure of the resulting topology $G[s]$ if the strategies s form a Nash equilibrium in I_k . First, it can be shown that in $G[s]$, two peers in the same cluster are always connected by a path that does not leave the cluster. In the absence of such an *intra-cluster path*, the shortest path between two peers in the same cluster has length at least $2 - 4\delta$ ($1 - 2\delta$ is the minimal inter-cluster distance in I_k). Hence, there is a stretch of at least $\frac{2-4\delta}{\epsilon/n}$ between this pair of peers. By constructing an intra-cluster link at cost α , a peer can significantly reduce these stretches, rendering such a link worthwhile. Furthermore, it can be shown that there exists exactly one link in both directions between clusters Π_a and Π_b , Π_b and Π_c , as well as between Π_1 and Π_2 . In all three cases, the argument is based on the fact that without such a link, the shortest path between peers in these neighboring clusters is via a third cluster, thus significantly increasing the stretches. Specifically, the sum of the stretches from a peer in one cluster to all peers in the neighboring cluster would exceed $\frac{k(4-\delta_{1a})}{1+\delta_{ab}}$. Because a single link to a peer in this neighboring cluster can reduce each stretch to roughly $1 + \epsilon$, the cost of connecting directly to a peer in the neighboring cluster is less than $\alpha + k(1 + \epsilon)$. Hence, if there is no link between two neighboring clusters, a peer has an incentive to construct such a link.

A third structural characteristic that can be derived for any Nash equilibrium is that for every i and j , there is *at most one* directed link from a cluster Π_i to peers in a cluster Π_j . Since ϵ is small and all peers are linked within their clusters, peer $\pi_i \in \Pi_i$ reduces its cost by dropping its link to cluster Π_j , if another peer in Π_i also has a link to a peer in Π_j .

To preserve connectivity, some peers in Π_1 and Π_2 must have links to top-peers. Based on the above observation that there is at most one link between two clusters in each direction, Lemma 5.2

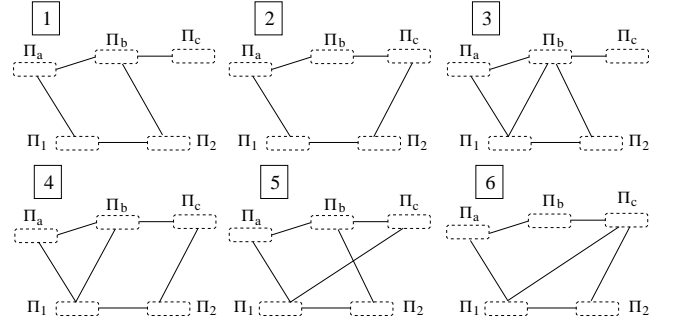


Figure 3: Candidates for a Nash equilibrium.

further narrows down the set of possible strategies for connecting peers in Π_1 and Π_2 to top-peers.

LEMMA 5.2. *In any Nash equilibrium of instance I_k , it holds that*

- i) *Neither peers in Π_1 nor Π_2 select three links to top-peers.*
- ii) *There exists a peer $\pi_1 \in \Pi_1$ that establishes a link to Π_a .*
- iii) *There is exactly one link from cluster Π_2 to either cluster Π_b or Π_c , but there is no link to Π_a .*

Correctness of all three properties is proven by verifying that there exists some node $\pi_1 \in \Pi_1$ or $\pi_2 \in \Pi_2$ that has an incentive to change its strategy in case the property is not satisfied. If, for instance, there are two peers $\pi_2, \pi_2' \in \Pi_2$ that simultaneously maintain links to Π_b and Π_c , (e.g. π_2 to Π_b and π_2' to Π_c , thus violating case iii), π_2' can lower its costs by dropping its link to Π_c . Intuitively, this holds because the sum of the stretches $\sum_{\pi_c \in \Pi_c} \text{stretch}(\pi_2', \pi_c)$ entailed by the indirection $\pi_2' \rightsquigarrow \pi_2 \rightsquigarrow \Pi_b \rightsquigarrow \Pi_c$ does not justify the additional cost α .

It can be shown that only the six structures depicted in Figure 3 remain valid candidates for Nash topologies. In each scenario, however, at least one peer benefits from deviating from its current strategy.

Case 1: In this case, a peer $\pi_1 \in \Pi_1$ can reduce its cost by adding a link ℓ_{1b} to a peer in Π_b .

Case 2: If the only outgoing link from Π_1 to a top-cluster is to cluster Π_a , the peer $\pi_2 \in \Pi_2$ maintaining the link to Π_c can be shown to profit from switching its link from Π_c to Π_b .

Case 3: The availability of ℓ_{1b} changes the optimal choice of the above mentioned peer $\pi_2 \in \Pi_2$. Unlike in the previous case, π_2 now prefers linking to Π_c instead of Π_b .

Case 4: Due to the existence of a link from a peer $\pi_2 \in \Pi_2$ to Π_c , the peer $\pi_1 \in \Pi_1$ with the link to Π_b has an incentive to drop this link ℓ_{1b} and instead use the detours via Π_2 and Π_a to connect to Π_c and Π_b , respectively.

Case 5: In this case, the peer $\pi_1 \in \Pi_1$ having the link to Π_c reduces its cost by replacing this link with a link to a peer in Π_b .

Case 6: Finally, this case is similar to Case 4 in the sense that $\pi_1 \in \Pi_1$ with the link to Π_b has an incentive to remove ℓ_{1c} .

This proof highlights how the system is ultimately trapped in an infinite loop of strategy changes, without ever converging to a stable situation. There is always at least one peer which can reduce its cost by changing its strategy. For instance, the following sequence of topology changes could repeat forever (cf. Figure 3): $1 \rightsquigarrow 3 \rightsquigarrow 4 \rightsquigarrow 2 \rightsquigarrow 1 \rightsquigarrow 3 \dots$. In other words, selfish peers will not achieve a stable network topology.

6. CONCLUSION

In order to gain a thorough understanding of the behavior of selfish peers, aspects *beyond* the contribution of resources have to be taken into account. In this paper, we have studied the impact of selfishness on the topology of a P2P network, that is, if utility-maximizing peers select their neighbors such that both latencies and link maintenance overhead are small. Our results show that topologies may degrade more severely when selfish peers value maintenance cost relatively higher than latency costs. Moreover, we prove that even in the absence of mobility or churn, a P2P network consisting of selfish peers may never reach a stable state. We consider this paper as a first step, and hope it sparks further research in this area. It would be interesting, for instance, to incorporate aspects such as overlay routing and congestion into our model.

ACKNOWLEDGMENTS

We are indebted to Uri Nadav for pointing us to [8]. We would also like to thank Yvonne Anne Oswald for many fruitful discussions.

REFERENCES

- [1] I. Abraham, A. Badola, D. Bickson, Dahlia Malkhi, Sharad Maloo, and Saar Ron. Practical Locality-Awareness for Large Scale Information Sharing. In *IPTPS*, 2005.
- [2] Ittai Abraham, Dahlia Malkhi, and Oren Dobzinski. LAND: Stretch $(1 + \epsilon)$ Locality Aware Networks for DHTs. In *SODA*, 2004.
- [3] E. Adar and B. Huberman. Free Riding on Gnutella. *First Monday*, 5(10), 2000.
- [4] Susanne Albers, Stefan Eilts, Eyal Even-Dar, Yishay Mansour, and Liam Roditty. On Nash Equilibria for a Network Creation Game. In *SODA*, 2006.
- [5] Nazareno Andrade, Miranda Mowbray, Aliandro Lima, Gustavo Wagner, and Matei Ripeanu. Influences on Cooperation in BitTorrent Communities. In *SIGCOMM*, 2005.
- [6] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *FOCS*, 2004.
- [7] J. Corbo and D. C. Parkes. The Price of Selfish Behavior in Bilateral Network Formation. In *PODC*, 2005.
- [8] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a Network Creation Game. In *PODC*, 2003.
- [9] Seung Jun and Mustaque Ahamad. Incentives in BitTorrent Induce Free Riding. In *SIGCOMM*, 2005.
- [10] C. Plaxton, R. Rajaraman, and A. W. Richa. Accessing Nearby Copies of Replicated Objects in a Distributed Environment. In *SPAA*, 1997.
- [11] A. Rowstron and P. Druschel. Pastry: Scalable, Distributed Object Location and Routing for Large-scale Peer-to-Peer Systems. In *IFIP/ACM Int. Conf. Dist. Sys. Platforms (Middleware)*, 2001.
- [12] B. Wong, A. Slivkins, and E. G. Sirer. Meridian: A Lightweight Network Location Service without Virtual Coordinates. In *SIGCOMM*, 2005.
- [13] M. Yang, Z. Zhang, X. Li, and Y. Dai. An Empirical Study of Free-Riding Behavior in the Maze P2P File Sharing System. In *IPTPS*, 2005.
- [14] Ben Y. Zhao, Ling Huang, Jeremy Stribling, Anthony D. Joseph, and John D. Kubiatowicz. Tapestry: A Resilient Global-scale Overlay for Service Deployment. *IEEE J. Selected Areas in Comm.*, 22(1), 2004.